COMBINED MODELING OF SPARSE AND DENSE NOISE FOR IMPROVEMENT OF RELEVANCE VECTOR MACHINE

Martin Sundin, Saikat Chatterjee and Magnus Jansson

ACCESS Linnaeus Center, School of Electrical Engineering
KTH Royal Institute of Technology, Stockholm, Sweden
masundi@kth.se, sach@kth.se, janssonm@kth.se

arXiv:1501.02579v1 [stat.ML] 12 Jan 2015

ABSTRACT

Using a Bayesian approach, we consider the problem of recovering sparse signals under additive sparse and dense noise. Typically, sparse noise models outliers, impulse bursts or data loss. To handle sparse noise, existing methods simultaneously estimate the sparse signal of interest and the sparse noise of no interest. For estimating the sparse signal, without the need of estimating the sparse noise, we construct a robust Relevance Vector Machine (RVM). In the RVM, sparse noise and ever present dense noise are treated through a combined noise model. The precision of combined noise is modeled by a diagonal matrix. We show that the new RVM update equations correspond to a non-symmetric sparsity inducing cost function. Further, the combined modeling is found to be computationally more efficient. We also extend the method to block-sparse signals and noise with known and unknown block structures. Through simulations, we show the performance and computation efficiency of the new RVM in several applications: recovery of sparse and block sparse signals, housing price prediction and image denoising.

1. INTRODUCTION

Noise modeling has an important role in the Bayesian inference setup to achieve better robustness and accuracy. Typically noise is considered to be additive and dense (possibly even white) in nature. In this paper we investigate the effect of sparse noise modeling in a standard Bayesian inference tool called the *Relevance Vector Machine* (RVM) [1].

The RVM is a Bayesian sparse kernel technique for applications in regression and classification [1]. Interest in the RVM can be attributed to the cause that it shares many characteristics of the popular support vector machine whilst providing Bayesian advantages [2, 3, 4], mainly providing posteriors for the object of interest. Generally the RVM is a fully Bayesian technique that aims to learn all the relevant system parameters iteratively to infer the object of interest. In a linear model setup used for regression, RVM introduces sparsity through a weight vector where the weights are essential to form linear combinations of relevant kernels to predict the object of interest; the weight vector is a set of system parameters and its sparsity leads to reduction of model complexity for regression. Naturally, the RVM has been further used for sparse representation techniques as well as developing Bayesian compressive sensing methods [5].

For a Bayesian linear model, the standard RVM uses a multivariate isotropic Gaussian prior to model the additive dense noise. Here isotropic means that the associated covariance matrix is proportional to the identity matrix. Such a dense noise model has inherent limitations to accommodate instances of outliers [6, 7, 8, 9, 10], impulse bursts [11, 12] or missing (lost) data [13, 14]. We hypothesize that a sparse and dense noise model can accommodate for the statistics of a variety of noise types, without causing degradation in performance for any noise type compared to the standard case of using only a dense noise model. In this paper, we develop RVM for such a combined (joint) sparse and dense noise scenario.

1.1. System model

We consider the following linear system model

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{e} + \mathbf{n}, \tag{1}$$

where $\mathbf{y} \in \mathbb{R}^m$ is the measurements, $\mathbf{x} \in \mathbb{R}^n$ is a sparse vector (for example weights in regression or sparse signal to estimate in compressed sensing), $\mathbf{A} \in \mathbb{R}^{m \times n}$ is a known system matrix (for example, regressors or sampling system). Further, $\mathbf{e} \in \mathbb{R}^m$ is sparse noise and $\mathbf{n} \in \mathbb{R}^m$ is dense noise. Using ℓ_0-norm notation to represent the number of non-zeros in a vector, we assume that $||\mathbf{x}||_0 \ll n$ and $||\mathbf{e}||_0 \ll m$ are small and unknown. The random vectors $\mathbf{x}$, $\mathbf{e}$ and $\mathbf{n}$ are independent. The model (1) is used in face recognition [15], image denoising [6] and compressed sensing [5].

1.2. Our contribution

We develop a RVM for the model (1), by treating $\mathbf{e} + \mathbf{n}$ as a combined noise. By learning parameters of $\mathbf{x}$ and $\mathbf{e} + \mathbf{n}$,

This work was partially supported by the Swedish Research Council under contract 621-2011-5847.

we estimate $\mathbf{x}$ without the need of estimating $\mathbf{e}$. We also consider the scenario where the signal $\mathbf{x}$ and noise $\mathbf{e}$ are block sparse. By using techniques similar to the ones in [16] we generalize the methods to signals with unknown block structure. The main technical contribution is to derive update equations that are used iteratively for estimation of parameters in the new RVM. We refer to the new RVM as the RVM for combined sparse and dense noise (SD-RVM). By an approximate analysis, the SD-RVM algorithm is shown to be equivalent to the minimization of a non-symmetric sparsity inducing cost function. Finally, the performance of SD-RVM is evaluated numerically using examples from compressed sensing, block sparse signal recovery, house price prediction and image denoising. Throughout the paper, we take an approach of comparing SD-RVM vis-a-vis the existing Robust Bayesian RVM (RB-RVM) [6] (described in the next section).

1.3. Prior work

To establish relevance of our work we briefly describe prior work in this section. Almost all prior works [6, 7, 8, 9] translate the linear setup (1) to the equivalent setup

$$\mathbf{y} = \begin{bmatrix} \mathbf{A} & \mathbf{I}_m \end{bmatrix} \begin{bmatrix} \mathbf{x} \\ \mathbf{e} \end{bmatrix} + \mathbf{n}, \tag{2}$$

where $\mathbf{I}_m$ is the $m \times m$ identity matrix, $\begin{bmatrix} \mathbf{A} & \mathbf{I}_m \end{bmatrix}$ acts as the effective system matrix and $\begin{bmatrix} \mathbf{x}^\top & \mathbf{e}^\top \end{bmatrix}^\top$ acts as the parameter vector to be estimated. The RB-RVM of [6] uses the standard RVM approach for (2) directly. Hence RB-RVM learns model parameters for all three signals $\mathbf{x} = [x_1, x_2, \ldots, x_n]^\top$, $\mathbf{e} = [e_1, e_2, \ldots, e_m]^\top$ and $\mathbf{n}$, and thus estimates both $\mathbf{x}$ and $\mathbf{e}$ jointly. RB-RVM assumes Gaussian priors

$$\mathbf{x} \sim \prod_{i=1}^{n} \mathcal{N}(0, \gamma_i^{-1}),\ \mathbf{e} \sim \prod_{i=1}^{m} \mathcal{N}(0, \nu_i^{-1}),\ \mathbf{n} \sim \mathcal{N}(\mathbf{0}, \beta^{-1}\mathbf{I}_m),$$

where the precisions (inverse variances) α_i, ν_i and β are unknown. The precisions are given Gamma priors

$$\begin{aligned} p(\gamma_i) &= \text{Gamma}(\gamma_i | a+1, b), \\ p(\nu_i) &= \text{Gamma}(\nu_i | a+1, b), \\ p(\beta) &= \text{Gamma}(\beta | c+1, d), \end{aligned} \tag{3}$$

where $\text{Gamma}(\gamma_i | a+1, b) \propto \gamma_i^a e^{-b\gamma_i}$ [1]. Typical practice is to maximize $p(\mathbf{y}|\boldsymbol{\gamma}, \boldsymbol{\nu}, \beta)$ to infer the precisions, where we used boldface symbols to denote the vectors

$$\begin{aligned} \boldsymbol{\gamma} &= [\gamma_1, \gamma_2, \ldots, \gamma_n]^\top, \\ \boldsymbol{\nu} &= [\nu_1, \nu_2, \ldots, \nu_m]^\top. \end{aligned}$$

Instead we take the alternative (full Bayesian) approach of maximizing $p(\mathbf{y}, \boldsymbol{\gamma}, \boldsymbol{\nu}, \beta)$ and assume that precisions have non-informative prior by taking the limit $(a, b, c, d) \to \mathbf{0}$. For the distributions considered here, maximizing the conditional distribution $p(\mathbf{y}|\boldsymbol{\gamma}, \boldsymbol{\nu}, \beta)$ becomes equivalent to maximizing the joint distribution $p(\mathbf{y}, \boldsymbol{\gamma}, \boldsymbol{\nu}, \beta)$, in the limit of non-informative priors. In calculations, however, the parameters (a, b, c, d) are often given small values to avoid numerical instabilities. To estimate $\begin{bmatrix} \mathbf{x}^\top & \mathbf{e}^\top \end{bmatrix}^\top$, RB-RVM fixes the precisions and sets

$$\begin{aligned} \begin{bmatrix} \hat{\mathbf{x}}^\top & \hat{\mathbf{e}}^\top \end{bmatrix}^\top &= \beta \boldsymbol{\Sigma}_{RB} [\mathbf{A}\ \mathbf{I}_m]^\top \mathbf{y}, \\ \boldsymbol{\Sigma}_{RB} &= \left(\begin{pmatrix} \boldsymbol{\Gamma} & \mathbf{0} \\ \mathbf{0} & \mathbf{N} \end{pmatrix} + \beta [\mathbf{A}\ \mathbf{I}_m]^\top [\mathbf{A}\ \mathbf{I}_m] \right)^{-1}, \end{aligned} \tag{4}$$

where $\boldsymbol{\Gamma} = \text{diag}(\gamma_1, \gamma_2, \ldots, \gamma_n)$ and $\mathbf{N} = \text{diag}(\nu_1, \nu_2, \ldots, \nu_m)$. The RB-RVM iteratively updates the precisions by maximizing $p(\mathbf{y}, \boldsymbol{\gamma}, \boldsymbol{\nu}, \beta)$, resulting in the update equations

$$\begin{aligned} \gamma_i^{new} &= \frac{1 - \gamma_i [\boldsymbol{\Sigma}_{RB}]_{ii}}{\hat{x}_i^2},\ \ \nu_i^{new} = \frac{1 - \nu_i [\boldsymbol{\Sigma}_{RB}]_{n+i,n+i}}{\hat{e}_i^2}, \\ \beta^{new} &= \frac{\sum_{i=1}^{n} \gamma_i [\boldsymbol{\Sigma}_{RB}]_{ii} + \sum_{j=1}^{m} \nu_j [\boldsymbol{\Sigma}_{RB}]_{n+i,n+i}}{||\mathbf{y} - \mathbf{A}\hat{\mathbf{x}} - \hat{\mathbf{e}}||_2^2}, \end{aligned} \tag{5}$$

where $[\boldsymbol{\Sigma}_{RB}]_{ii}$ denotes the (i, i) component of the matrix $\boldsymbol{\Sigma}_{RB}$.

The update equations (4) and (5) are found by applying the standard RVM to (2). Derivations can be found in e.g. [1, 2]. Iterating until convergence gives the final estimates $\hat{\mathbf{x}}$ and $\hat{\mathbf{e}}$. In the iterations, some precisions become large, making their respective components in $\hat{\mathbf{x}}$ and $\hat{\mathbf{e}}$ close to zero. This makes the final estimate of $\hat{\mathbf{x}}$ and $\hat{\mathbf{e}}$ sparse.

RVM has high similarity with Sparse Bayesian Learning (SBL) [16, 17, 3, 4]. Sparse Bayesian learning has been used for structured sparse signals, for example block sparse signals [16], where the problem of unknown signal block structure was treated using overlapping blocks. The model extension of RB-RVM shown in (2) for handling block sparse noise with unknown block structure is straight-forward to derive. However, in our formulation, as we are not estimating the noise explicitly, the use of block sparse noise with unknown block structure is non-trivial.

Further, non Bayesian (even not statistical) methods have been used for sparse estimation problems [7, 18, 19, 20]. For example, the ℓ_1-norm minimization method justice pursuit (JP) [7] uses the optimization technique of the standard basis pursuit denoising method [18], as follows

$$\hat{\mathbf{x}}, \hat{\mathbf{e}} = \arg\min_{\mathbf{x},\mathbf{e}} ||\mathbf{x}||_1 + ||\mathbf{e}||_1 \text{ s.t. } ||\mathbf{y} - \mathbf{A}\mathbf{x} - \mathbf{e}||_2 \leq \epsilon, \tag{6}$$

where $\epsilon > 0$ is a model parameter set by the user. For unknown noise power, it is impossible to know ϵ a-priori. We mention that a fully Bayesian setup like the RVM does not require parameters set by a user.

2. RVM FOR COMBINED SPARSE AND DENSE NOISE (SD-RVM)

2.1. SD-RVM Method

For (1), we propose to use a combined model for the two additive noises, as follows

$$\mathbf{e}+\mathbf{n} \sim \mathcal{N}(\mathbf{0}, \mathbf{B}^{-1}), \tag{7}$$

where $\mathbf{B} = \mathrm{diag}(\beta_1, \beta_2, \ldots, \beta_m)$. We also use $\boldsymbol{\beta}$ to denote the vector $\boldsymbol{\beta} = [\beta_1, \beta_2, \ldots, \beta_m]^\top$. That means the two noises are treated as a single combined noise where each noise component has its own precision. The rationale is that we do not need to seperate the two noises. Although our model promotes sparsity in the noise we empirically find that it is able to model sparse and non-sparse noise. Using the noise model (7) and that $\mathbf{x} \sim \prod_{i=1}^n \mathcal{N}(0, \gamma_i^{-1})$, we find the maximum a posteriori (MAP) estimate

$$\begin{aligned}\hat{\mathbf{x}} &= \boldsymbol{\Sigma}\mathbf{A}^\top\mathbf{B}\mathbf{y},\\ \boldsymbol{\Sigma} &= (\boldsymbol{\Gamma} + \mathbf{A}^\top\mathbf{B}\mathbf{A})^{-1},\end{aligned}$$

where as before $\boldsymbol{\Gamma} = \mathrm{diag}(\gamma_1, \gamma_2, \ldots, \gamma_n)$. The precisions are updated as

$$\gamma_i^{new} = \frac{1-\gamma_i\Sigma_{ii}}{\hat{x}_i^2}, \tag{8}$$

$$\beta_j^{new} = \frac{1-\beta_j[\mathbf{A}\boldsymbol{\Sigma}\mathbf{A}^\top]_{jj}}{[\mathbf{y}-\mathbf{A}\hat{\mathbf{x}}]_j^2}, \tag{9}$$

where $\Sigma_{ii} = [\boldsymbol{\Sigma}]_{ii}$. The derivations of (8) and (9) are given in the next section.

2.2. Derivation of update equations for SD-RVM

To update the precisions we maximize the distribution $p(\mathbf{y}, \boldsymbol{\gamma}, \boldsymbol{\beta}) = p(\mathbf{y}|\boldsymbol{\gamma}, \boldsymbol{\beta})p(\boldsymbol{\gamma})p(\boldsymbol{\beta})$ (obtained by marginalizing over $\mathbf{x}$), with respect to γ_i and β_j, where we use the prior

$$p(\beta_j) = \mathrm{Gamma}(\beta_j|c+1, d),$$

and $p(\gamma_i)$ is as in (3). The log-likelihood of the parameters is

$$\begin{aligned}\mathcal{L} =& \text{constant} - \frac{1}{2}\log\det(\mathbf{B}^{-1} + \mathbf{A}\boldsymbol{\Gamma}^{-1}\mathbf{A}^\top) \\ &- \frac{1}{2}\mathbf{y}^\top(\mathbf{B}^{-1} + \mathbf{A}\boldsymbol{\Gamma}^{-1}\mathbf{A}^\top)^{-1}\mathbf{y} \\ &+ \sum_{i=1}^n(a\log\gamma_i - b\gamma_i) + \sum_{j=1}^m(c\log\beta_j - d\beta_j).\end{aligned} \tag{10}$$

We maximize $\mathcal{L}$ w.r.t. γ_i by setting the derivative to zero. To simplify the derivative we use that

$$\frac{\partial}{\partial\gamma_i}\left(\mathbf{y}^\top(\mathbf{B}^{-1} + \mathbf{A}\boldsymbol{\Gamma}^{-1}\mathbf{A}^\top)^{-1}\mathbf{y}\right) = \hat{x}_i^2, \tag{11}$$

and the determinant lemma [21]

$$\det(\mathbf{B}^{-1} + \mathbf{A}\boldsymbol{\Gamma}^{-1}\mathbf{A}^\top) = \det(\boldsymbol{\Sigma}^{-1})\det(\boldsymbol{\Gamma}^{-1})\det(\mathbf{B}^{-1}). \tag{12}$$

Using (11) and (12) we find that $\mathcal{L}$ is maximized w.r.t. γ_i when

$$-\frac{1}{2}\Sigma_{ii} + \frac{1}{2\gamma_i} + \frac{a}{\gamma_i} - b - \frac{1}{2}\hat{x}_i^2 = 0. \tag{13}$$

Instead of solving for γ_i (which would require solving a non-linear coupled equation since $\boldsymbol{\Sigma}$ and $\hat{\mathbf{x}}$ depend on γ_i) we approximate the equation as

$$1 - \gamma_i\Sigma_{ii} + 2a - (\hat{x}_i^2 + 2b)\gamma_i^{new} = 0. \tag{14}$$

We solve (14) for γ_i^{new} rather than (32) for γ_i since it in practice often results in a better convergence [1, 22]. The update equation then becomes

$$\gamma_i^{new} = \frac{1-\gamma_i\Sigma_{ii} + 2a}{\hat{x}_i^2 + 2b}.$$

Setting $a = b = 0$ we obtain (8).

For the noise precisions we use that

$$\frac{\partial}{\partial\beta_j}\left[\mathbf{y}^\top(\mathbf{B}^{-1} + \mathbf{A}\boldsymbol{\Gamma}^{-1}\mathbf{A}^\top)^{-1}\mathbf{y}\right] = [\mathbf{y}-\mathbf{A}\hat{\mathbf{x}}]_j^2. \tag{15}$$

We find that $\mathcal{L}$ is maximized w.r.t. β_j when

$$-\frac{1}{2}\mathrm{tr}(\boldsymbol{\Sigma}\mathbf{A}_{j,:}^\top\mathbf{A}_{j,:}) + \frac{1}{2\beta_j} - \frac{1}{2}[\mathbf{y}-\mathbf{A}\hat{\mathbf{x}}]_j^2 + \frac{c}{\beta_j} - d = 0,$$

where $\mathbf{A}_{j,:}$ denotes the j'th row vector of $\mathbf{A}$. Rewriting the equation as

$$1 - \beta_j\mathbf{A}_{j,:}\boldsymbol{\Sigma}\mathbf{A}_{j,:}^\top + 2c - ([\mathbf{y}-\mathbf{A}\hat{\mathbf{x}}]_j^2 + 2d)\beta_j^{new} = 0,$$

using that $\mathbf{A}_{j,:}\boldsymbol{\Sigma}\mathbf{A}_{j,:}^\top = [\mathbf{A}\boldsymbol{\Sigma}\mathbf{A}^\top]_{jj}$, we find that

$$\beta_j^{new} = \frac{1-\beta_j[\mathbf{A}\boldsymbol{\Sigma}\mathbf{A}^\top]_{jj} + 2c}{[\mathbf{y}-\mathbf{A}\hat{\mathbf{x}}]_j^2 + 2d}.$$

Setting $c = d = 0$ we obtain (9).

The derivations of (11) and (15) are given in Appendix 6.1.

2.3. Analysis of sparsity

Several approximations are made in the derivation of the iterative update equations. It is interesting to see how the approximations affect the sparsity of the solution. In this subsection, we show that the approximations make the SD-RVM equivalent to minimizing a non-symmetric sparsity promoting cost function.

To motivate that the standard RVM is sparsity promoting, one can use that the marginal distribution of x_i is a student-t

distribution. For a fixed β (and $\mathbf{e} = \mathbf{0}$), the standard RVM is therefore an iterative method for solving (details can be found in [1])

$$\min_{\mathbf{x}} \frac{\beta}{2}||\mathbf{y} - \mathbf{A}\mathbf{x}||_2^2 + \left(1 + \frac{a}{2}\right) \sum_{i=1}^{n} \log(x_i^2 + 2b).$$

The log-sum cost function can be used as a sparsity promoting cost function, making it plausible that the RVM promotes sparsity.

For the SD-RVM, the precisions are updated by maximizing the marginal distribution $p(\mathbf{y}, \boldsymbol{\gamma}, \boldsymbol{\beta})$. The problem is equivalent to maximizing $\mathcal{L}$ in (31). We show approximations for relevant parts of the right hand side of $\mathcal{L}$ as follows

$$\log\det(\boldsymbol{\Sigma}^{-1}) \approx \log\det((\boldsymbol{\Sigma}^{old})^{-1}) + \sum_{i=1}^{n} \Sigma_{ii}^{old}(\gamma_i - \gamma_i^{old}) + \sum_{j=1}^{m} [\mathbf{A}\boldsymbol{\Sigma}^{old}\mathbf{A}^\top]_{jj}(\beta_j - \beta_j^{old}), \quad (16)$$

where the approximation is up to first order in $\boldsymbol{\gamma}$ and $\boldsymbol{\beta}$. We rewrite the problem in variables $\mathbf{x}$ and $\tilde{\mathbf{e}}$ using that [23]

$$\mathbf{y}^\top(\mathbf{A}\boldsymbol{\Gamma}^{-1}\mathbf{A}^\top + \mathbf{B}^{-1})^{-1}\mathbf{y} = \min_{\mathbf{x},\tilde{\mathbf{e}}} \sum_{i=1}^{n} \gamma_i x_i^2 + \sum_{j=1}^{m} \beta_j \tilde{e}_j^2, \quad \text{such that } \mathbf{A}\mathbf{x} + \tilde{\mathbf{e}} = \mathbf{y} \quad (17)$$

where now $\tilde{\mathbf{e}} = \mathbf{e} + \mathbf{n}$ as in (7). Under the approximation (16) and the reformulation (17), maximization of $\log p(\mathbf{y}, \boldsymbol{\gamma}, \boldsymbol{\beta})$ becomes equivalent to

$$\min_{\gamma_i,\beta_j,\mathbf{x},\tilde{\mathbf{e}}} \sum_{i=1}^{n} \left[(x_i^2 + \Sigma_{ii}^{old} + 2b)\gamma_i + (1 + 2a)\log(\gamma_i)\right] + \sum_{j=1}^{m} \left[(e_j^2 + [\mathbf{A}\boldsymbol{\Sigma}^{old}\mathbf{A}^\top]_{jj} + 2d)\beta_j + (1 + 2c)\log(\beta_j)\right].$$
$$\text{such that } \mathbf{A}\mathbf{x} + \tilde{\mathbf{e}} = \mathbf{y}$$

By minimizing the objective with respect to γ_i and β_j, the problem reduces to

$$\min_{\mathbf{x},\tilde{\mathbf{e}}} (1 + 2a) \sum_{i=1}^{n} \log(x_i^2 + \Sigma_{ii}^{old} + 2b) + (1 + 2c) \sum_{j=1}^{m} \log(\tilde{e}_j^2 + [\mathbf{A}\boldsymbol{\Sigma}^{old}\mathbf{A}^\top]_{jj} + 2d), \quad (18)$$
$$\text{such that } \mathbf{A}\mathbf{x} + \tilde{\mathbf{e}} = \mathbf{y}$$

where we have ignored additive constants. Because of the approximations, the constants Σ_{ii}^{old} and $[\mathbf{A}\boldsymbol{\Sigma}^{old}\mathbf{A}^\top]_{jj}$ make the cost function non-symmetric in the components of $\mathbf{x}$ and $\tilde{\mathbf{e}}$. The SD-RVM is thus equivalent to minimizing a non-symmetric sparsity promoting cost function. In a similar way it can be shown that the standard RVM and RB-RVM are also equivalent to non-symmetric cost functions under appropriate approximations. A two-dimensional example using $\mathbf{x} = [x_1, x_2]^\top$ is shown in Fig. 1.

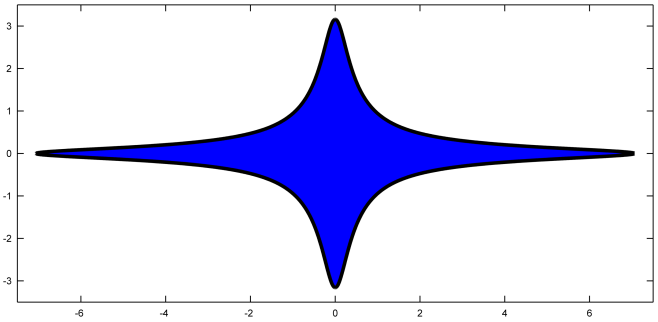

Fig. 1. The non-symmetric *log-ball* $\log(x_1^2 + 0.02) + \log(x_2^2 + 0.1) \leq 0$. SD-RVM is equivalent to finding the smallest non-symmetric log-ball that intersects the linear subspace $\mathbf{A}\mathbf{x} + \tilde{\mathbf{e}} = \mathbf{y}$.

2.4. Computational complexity

In this section we take a non-rigorous approach for quantifying the computational complexity of SD-RVM. The complexity is computed per iteration, since the number of iterations depends on the stopping criterion used, and with the assumption of a naive implementation. Each iteration of SD-RVM requires $\mathcal{O}(n^3)$ flops to compute the matrix $\boldsymbol{\Sigma}$ using Gauss-Jordan elimination [24]. Updating the precisions requires $\mathcal{O}(nm)$ flops since the residual $\mathbf{y} - \mathbf{A}\hat{\mathbf{x}}$ needs to be computed. Hence the computational complexity of SD-RVM is

$$\mathcal{O}(\max(nm, n^3)) = \mathcal{O}(n \cdot \max(m, n^2)).$$

A natural interest is the complexity of RB-RVM. Again with the assumption of a naive implementation, each iteration of RB-RVM requires the inversion of a $(n+m)\times(n+m)$ matrix to compute $\boldsymbol{\Sigma}_{RB}$. Updating the precisions requires $\mathcal{O}(nm)$ flops and hence the computational complexity of RB-RVM is

$$\mathcal{O}(\max(nm, (n+m)^3)) = \mathcal{O}((n+m)^3).$$

In Section 4.1 we provide numerical evaluations to quantify algorithm run time requirements that confirm that SD-RVM is typically faster than RB-RVM.

3. SD-RVM WITH BLOCK STRUCTURE

3.1. SD-RVM for known block structure

To describe a block sparse signal $\mathbf{x} \in \mathbb{R}^n$ with known block structure we partition $[n] = \{1, 2, \ldots, n\}$ into blocks as

$$[n] = I_1 \cup I_2 \cup \cdots \cup I_p,$$

where $|I_i| = n_i$ and $I_i \cap I_j = \emptyset$ for $i \neq j$. The signal is block sparse when only a few blocks of the signal are non-zero. The component-wise SD-RVM generalizes to this scenario by requiring that the precisions are equal in each block, i.e. we choose the prior distribution for the components of block I_i to be

$$\mathbf{x}_{I_i} \sim \mathcal{N}(\mathbf{0}, \gamma_i^{-1}\mathbf{I}_{n_i}).$$

where $\mathbf{x}_{I_i} \in \mathbb{R}^{n_i}$ denotes the vector consisting of the components of $\mathbf{x}$ with indices in I_i.

Similarly we can partition the components of the sparse noise $\tilde{\mathbf{e}} \in \mathbb{R}^m$ into blocks as

$$[m] = J_1 \cup J_2 \cup \cdots \cup J_q,$$

where $|J_j| = m_j$, $J_j \cap J_i = \emptyset$ for $i \neq j$ and the block J_j of $\mathbf{e}$ is given the prior distribution

$$\tilde{\mathbf{e}}_{J_j} \sim \mathcal{N}(\mathbf{0}, \beta_j^{-1}\mathbf{I}_{m_j}).$$

As before, the precisions are given gamma distributions (3) as priors, where now

$$p(\beta_j) = \text{Gamma}(\beta_j | c+1, d). \tag{19}$$

Using this model, we derive the update equations of precisions as below

$$\gamma_i^{new} = \frac{n_i - \gamma_i \operatorname{tr}(\mathbf{\Sigma}_{I_i}) + 2a}{||\hat{\mathbf{x}}_{I_i}||_2^2 + 2b}, \tag{20}$$

$$\beta_j^{new} = \frac{m_j - \beta_j \operatorname{tr}([\mathbf{A}\mathbf{\Sigma}\mathbf{A}^\top]_{J_j}) + 2c}{||(\mathbf{y} - \mathbf{A}\hat{\mathbf{x}})_{J_j}||_2^2 + 2d}, \tag{21}$$

where $\mathbf{\Sigma}_{I_i}$ denotes the $n_i \times n_i$ submatrix of $\mathbf{\Sigma}$ formed by elements appropriately indexed by I_i. By setting $I_i = \{i\}$ and $J_j = \{j\}$ we obtain the update equations for componentwise sparse signal and noise. We see that when $I_i = \{i\}$ and $J_j = J = [m]$, then (21) reduces to the update equations of the standard RVM since

$$n - \beta \operatorname{tr}(\mathbf{A}\mathbf{\Sigma}\mathbf{A}^\top) = \sum_i \gamma_i \Sigma_{ii}.$$

The derivation of the update equations (20) and (21) is found in Appendix 6.2.

3.2. SD-RVM for unknown block structure

In some situations the signal can have an unknown block structure, i.e. the signal is block sparse, but the dimensions and positions of the blocks are unknown. This scenario can be handled by treating the signal as a superposition of block sparse signals [16] (see illustration in Figure 2). This approach also describes the scenario (1) when $\mathbf{e}$ is component wise sparse and $\mathbf{n}$ is dense (e.g. Gaussian). The precision of each component is then a combination of the precisions

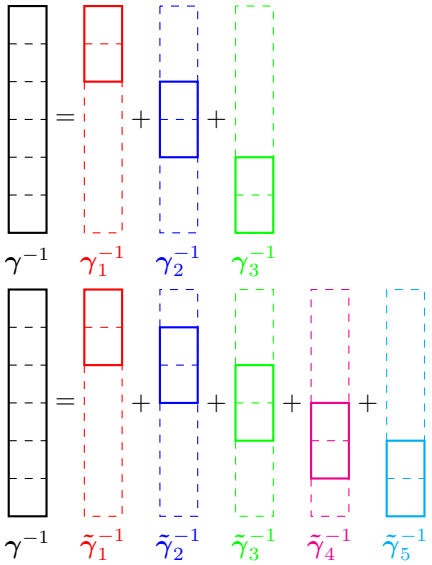

Fig. 2. Illustration of non-overlapping and overlapping block parameterizations.

of the blocks to which the component belongs. Let γ_i be the precision of the component x_i and $\tilde{\gamma}_k$ be the precision of block I_k. We model the signal as

$$\begin{aligned} x_i &\sim \mathcal{N}(0, \gamma_i^{-1}), \\ \gamma_i^{-1} &= \sum_{k, i \in I_k} \tilde{\gamma}_k^{-1}. \end{aligned} \tag{22}$$

We model the noise in a similar way with precisions β_j for component j and precisions $\tilde{\beta}_l$ for the block with support J_l. To promote sparsity, the precisions of the underlying blocks are given gamma distributions as priors

$$\begin{aligned} \tilde{\gamma}_k &\sim \text{Gamma}(\tilde{\gamma}_k | a+1, b), \\ \tilde{\beta}_l &\sim \text{Gamma}(\tilde{\beta}_l | c+1, d). \end{aligned}$$

In each iteration we update the underlying precisions $\tilde{\gamma}_k$. The componentwise precisions are then updated using (22). With this model, the update equations for the precisions become

$$\tilde{\gamma}_k^{new} = \frac{\frac{1}{\tilde{\gamma}_k}\operatorname{tr}(\mathbf{\Gamma}_k) - \frac{1}{\tilde{\gamma}_k}\operatorname{tr}(\mathbf{\Gamma}_k\mathbf{\Sigma}\mathbf{\Gamma}_k) + 2a}{\frac{1}{\tilde{\gamma}_k^2}||\mathbf{\Gamma}_i\hat{\mathbf{x}}||_2^2 + 2b}, \tag{23}$$

$$\tilde{\beta}_l^{new} = \frac{\frac{1}{\tilde{\beta}_l}\operatorname{tr}(\mathbf{B}_l) - \frac{1}{\tilde{\beta}_l}\operatorname{tr}(\mathbf{B}_l\mathbf{A}^\top\mathbf{\Sigma}\mathbf{A}\mathbf{B}_l) + 2c}{\frac{1}{\tilde{\beta}_l^2}||\mathbf{B}_l(\mathbf{y} - \mathbf{A}\hat{\mathbf{x}})||_2^2 + 2d}, \tag{24}$$

where $\mathbf{\Gamma}_k$ is the diagonal matrix with $[\mathbf{\Gamma}_k]_{ii} = \gamma_i$ if $i \in I_k$ and $[\mathbf{\Gamma}_k]_{ii} = 0$ otherwise. We denote the corresponding matrix for β_l by $\mathbf{B}_l$. The componentwise precisions are updated using (22) and similar for β_j.

We see that when the underlying blocks are disjoint, then $\gamma_i = \tilde{\gamma}_k$ for all $i \in I_k$ and $\beta_j = \tilde{\beta}_l$ for all $j \in J_l$. The

update equations then reduce to the update equations (21) for the block sparse model with known block structure.

3.2.1. *Sparse and dense noise*

In the model (1) where $\mathbf{x}$ and $\mathbf{e}$ are componentwise sparse and $\mathbf{n}$ is dense, then

$$I_i = \{i\},\ J_j = \{j\},\ J_{m+1} = [m], \tag{25}$$

where $i = 1, 2, \ldots, n$ and $j = 1, 2, \ldots, m$. In this scenario the support set of the sparse and dense noise is overlapping, so the update equations for the precisions become

$$\gamma_i^{new} = \frac{1 - \Sigma_{ii}\gamma_i + 2a}{\hat{x}_i^2 + 2b},$$

$$\tilde{\beta}_j^{new} = \frac{1 - \frac{\beta_j}{\tilde{\beta}_j}[\mathbf{A}^\top \mathbf{\Sigma} \mathbf{A}]_{jj} + 2c}{\beta_j [\mathbf{y} - \mathbf{A}\hat{\mathbf{x}}]_j^2 + 2d},\ j = 1, 2, \ldots, m,$$

$$\tilde{\beta}_{m+1}^{new} = \frac{\sum_{j=1}^m \beta_j - \frac{1}{\tilde{\beta}_{m+1}} \sum_{j=1}^m \beta_j^2 [\mathbf{A}^\top \mathbf{\Sigma} \mathbf{A}]_{jj} + 2c}{\sum_{j=1}^m \beta_j^2 [\mathbf{y} - \mathbf{A}\hat{\mathbf{x}}]_j^2 + 2d},$$

$$\beta_j = (\tilde{\beta}_j^{-1} + \beta_{m+1}^{-1})^{-1}.$$

We will use these update equations in the simulations where the signal is component-wise sparse and the noise is a sum of (component-wise) sparse and dense noise. It turns out that this method is slightly better than the SD-RVM in section 2.1.

4. SIMULATION EXPERIMENTS

In this section we evaluate the performance of the SD-RVM using several scenarios – for simulated and real signals. For simulated signals, we considered the sparse and block sparse recovery problem in compressed sensing. Then for real signals, we considered prediction of house prices using the Boston housing dataset [25] and denoising of images contaminated by *salt and pepper* noise. In the simulations we used the cvx toolbox [26] to implement JP.

4.1. Compressed sensing

The recovery problem in compressed sensing consists of estimating a sparse vector $\mathbf{x} \in \mathbb{R}^n$ in (1) from m linear measurements, where $m \ll n$. To evaluate the performance of the algorithms, we generated measurement matrices $\mathbf{A} \in \mathbb{R}^{m \times n}$ by drawing their components from a $\mathcal{N}(0, 1)$ distribution and scaling their column vectors to unit norm. We selected the positions of the active components of $\mathbf{x}$ and $\mathbf{e}$ uniformly at random and draw their values from $\mathcal{N}(0, 1)$. In the simulation we draw the additive noise $\mathbf{n}$ from $\mathcal{N}(\mathbf{0}, \sigma_n^2 \mathbf{I}_m)$. We compared JP, the standard RVM, RB-RVM and SD-RVM. For JP (6) we assumed σ_n to be known and set $\epsilon = \sigma_n \sqrt{m + 2\sqrt{2m}}$ as proposed in [27].

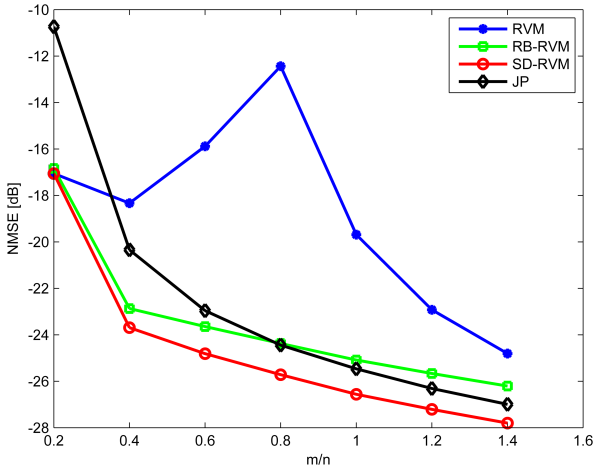

Fig. 3. NMSE vs. m/n for outlier free measurements.

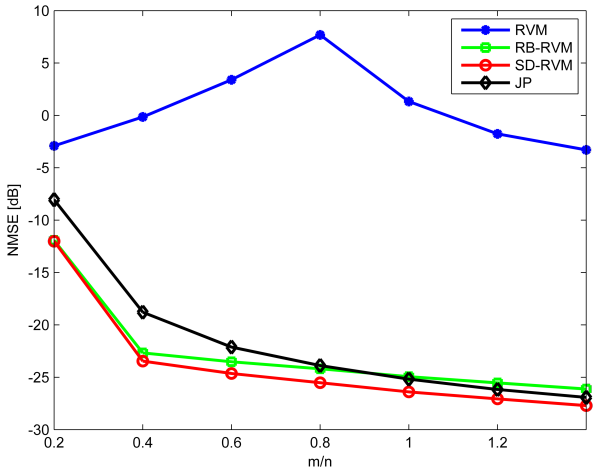

Fig. 4. NMSE vs. m/n for 5% outliers contaminated measurements.

In the simulations we varied the measurement rate m/n (ratio of the number of measurements and the signal dimension) for measurements without outliers and with 5% outliers. We chose $n = 100$ and fixed the signal-to-dense-noise-ratio (SDNR)

$$\text{SDNR} = E[||\mathbf{Ax}||_2^2]/E[||\mathbf{n}||_2^2] = ||\mathbf{x}||_0/(m\sigma_n^2),$$

to 20 dB. By generating 100 measurement matrices and 100 vectors $\mathbf{x}$ and $\mathbf{e}$ for each matrix we numerically evaluated the Normalized Mean Square Error (NMSE)

$$\text{NMSE} = E[||\mathbf{x} - \hat{\mathbf{x}}||_2^2]/E[||\mathbf{x}||_2^2].$$

The results are shown in Figure 3 and Figure 4. We found that SD-RVM outperformed the other methods. The improvement of SD-RVM over RB-RVM was 1 to 1.5 dB for $m/n > 0.5$, with and without outliers. Compared to JP, the improvement of SD-RVM was 3 to 3.7 without outlier noise and 1 to 4 dB with outlier noise when $m/n > 0.5$. The poor perfor-

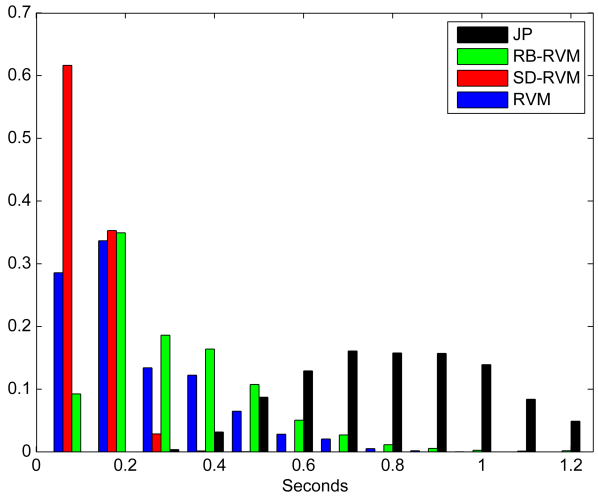

Fig. 5. Histogram of cputimes for the compressed sensing problem. Time is in seconds.

mance of RVM is due to sensitivity to the regularization parameters. The performance of RVM improves greatly when the regularization are optimally tuned, however, the optimal values varies with SNR and measurement dimensions. The experiments show that the performance of SD-RVM does not degrade in the absence of sparse noise.

For each realization of the problem we measured the runtime (cpu time) of each algorithm. The histogram of the runtimes is shown in figure 5. We found that the runtimes of the RVM algotithms (the standard RVM, RB-RVM and SD-RVM) were shorter than the runtime of JP and the runtimes of JP were spread over a larger range. Of the RVM algorithms, SD-RVM had the highest concentration of low runtime (≤ 0.2 seconds), while the runtimes of the standard RVM and RB-RVM was more concentrated around 0.2 seconds. The histogram in figure 5 has been truncated to only show percentage for the visible values .

4.2. Block sparse signals

The recovery problem in compressed sensing can be generalized to block sparse signals and noise [28]. For block sparse signals, the signal components are partitioned into blocks of which only a few blocks are non-zero. Sparse Bayesian learning (SBL) extended to the block sparse signal case is often referred to as block SBL (BSBL) [16, 17]. The problem of unknown block structure can be solved by overparametrizing the blocks [16]. In BSBL [16], the signal is modelled as

$$\mathbf{A}\mathbf{x} = \sum_{I \subset [n]} \mathbf{A}_I \mathbf{x}_I,$$

i.e. the measured signal is modelled as a sum of signals where each signal represents a block of the original signal.

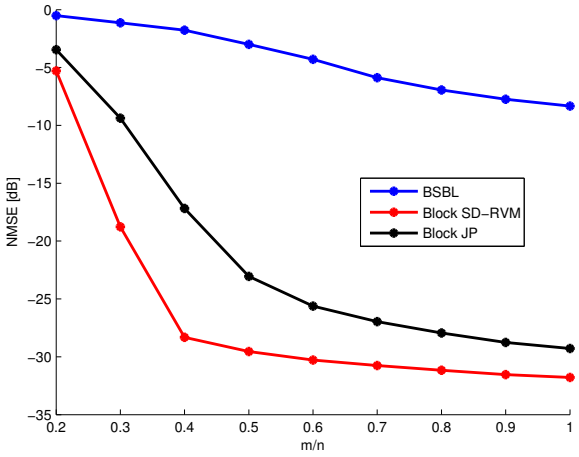

Fig. 6. NMSE vs. m/n for signals with known block structure and 5% outliers in measurements.

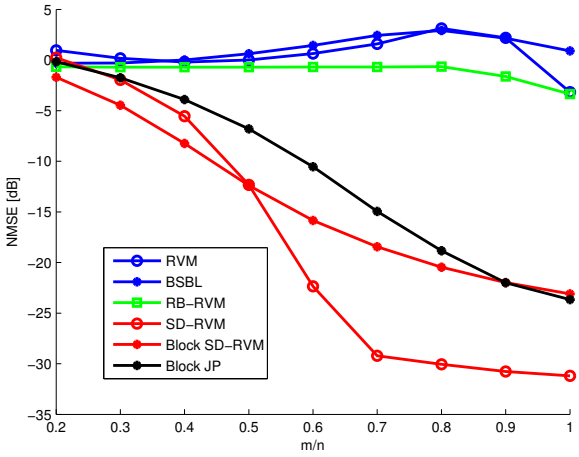

Fig. 7. NMSE vs. m/n for signals with unknown block structure and 5% outliers in measurements.

The resulting problem can then be solved using BSBL for known block structure. When the minimum block size K_{min} is known, the summation can be restricted to subsets of size $|I| = K_{min}$ [16].

The SD-RVM can be extended to the block sparse case using the methods developed in Section 3.1 and Section 3.2. Justice Pursuit can be extended to the block sparse case in a similar way as BSBL by setting

$$\hat{\mathbf{x}}, \hat{\mathbf{e}} = \arg\min_{\mathbf{x},\mathbf{e}} \sum_i ||\mathbf{x}_{I_i}||_2 + \sum_j ||\mathbf{e}_{J_j}||_2, \qquad (26)$$
$$\text{such that } ||\mathbf{y} - \mathbf{A}\mathbf{x} - \mathbf{e}||_2 \leq \epsilon$$

where the sum runs over all blocks (non-overlapping or overlapping) and as before we assume the noise variance to be known and set $\epsilon = \sigma_n \sqrt{m + \sqrt{8m}}$. For unknown block structure we also compared with component sparse methods RVM, RB-RVM and SD-RVM.

To numerically evaluate the performance of the block sparse algorithms we varied the measurement rate m/n for measurements with 5% sparse noise. We set the signal dimension to $n = 100$ and fixed the SDNR to 20 dB. We divided the signal $\mathbf{x}$ into 20 blocks of equal size of which 3 blocks were non-zero. The sparse noise consisted of blocks with 5 components in each block. In the sparse noise, 5% of the blocks were active. For known block structure, the blocks were choosen uniformly at random from a set of predefined and non-overlapping blocks while for unknown block structure, the first component of each block was choosen uniformly at random, making it possible for the blocks to overlap. The active components of the signal and the sparse noise were drawn from $\mathcal{N}(0, 1)$. By generating 50 measurement matrices $\mathbf{A}$ and 50 signals $\mathbf{x}$ and sparse noises $\mathbf{e}$ for each matrix we numerically evaluated the NMSE.

For known block structure we found that block SD-RVM outperformed the other methods. The NMSE of the block sparse SD-RVM was lower than the NMSE of block JP by more than 10 dB for $m/n = 0.3, 0.4$ and from 2 to 6 dB lower for $m/n \geq 0.5$. The results are presented in Figure 6.

For unknown block structure we found that for $m/n < 0.5$, SD-RVM for unknown block structure gave best performance while for $m/n \geq 0.5$, the usual component sparse SD-RVM gave the best perfomance. The NMSE of JP for unknown block structure was about 5 dB larger than the NMSE of block SD-RVM for $0.4 \leq m/n \leq 0.6$, while for $m/n \geq 0.9$ block JP gave a better NMSE than block SD-RVM. As expected, RVM and BSBL gave poor performance since they are not developed to handle measurements with sparse noise. The results are shown in figure 7.

4.3. House price prediction

One real world problem is the prediction of house prices. To test the algorithms on real data, we used the Boston housing dataset [25]. The dataset consists of 506 house prices in suburbs of Boston and 13 parameters (air quality, accessibility, pupil-to-teacher ratio, etc.) for each house. The problem is to predict the median house price for part of the dataset (test data) using the complement dataset (training data) to learn regression parameters. We model the house prices as

$$p_i = \mathbf{w}_i^\top \mathbf{x} + n_i + e_i,$$

where p_i is the price of house i, $\mathbf{w}_i \in \mathbb{R}^{13}$ contains the parameters of house i, $\mathbf{x} \in \mathbb{R}^{13}$ is the regression vector, n_i is (Gaussian) noise and e_i is a (possible) outlier. Very expensive or inexpensive houses can treated as outliers. The goal is to estimate the median house price for the test set. We find the median by estimating the regression parameters and setting

$$\hat{m} = \text{median}(\mathbf{W}^\top \hat{\mathbf{x}}),$$

where $\mathbf{W}$ contains the parameters of the houses in the test set. It is believed that only a few parameters are important to the average customer, $\mathbf{x}$ can therefore be modelled as a sparse vector.

Table 1. Prediction of median houseprice using the Boston Housing dataset. Mean error and mean cputime (in seconds) for different fractions, ρ, of the dataset used as training set.

	RVM		RB-RVM		SD-RVM	
ρ	**Error**	**Cputime**	**Error**	**Cputime**	**Error**	**Cputime**
0.3	1.24	0.18	0.43	0.60	0.38	0.15
0.4	1.26	0.29	0.39	1.25	0.35	0.25
0.5	1.27	0.42	0.39	2.20	0.36	0.38
0.6	1.28	0.60	0.41	3.28	0.37	0.53
0.7	1.28	0.92	0.45	5.27	0.43	0.80

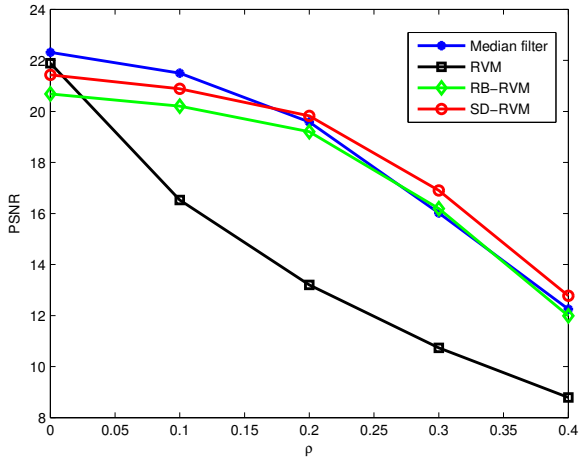

Fig. 8. PSNR vs. percentage of salt and pepper noise (ρ) averaged over 7 images with 10 noise realizations for each image for each value of ρ.

We used a fraction ρ of the dataset as training data and the rest as test set. By choosing the training set uniformly at random we evaluated the mean absolute error of the predicted median and mean cputime (in seconds) over 1000 realizations.

We found that SD-RVM gave 10% to 5% lower mean error than that of RB-RVM and the mean error of RB-RVM and SD-RVM was about 70% lower than the error of the RVM (see Table 1). The cputime of SD-RVM was 16% to 25% of the cputime of RB-RVM.

4.4. Image denoising

A grayscale image (represented in double-precision) can be modelled as an array of numbers in the interval from 0 to 1. Common sources of noise in images are electronic noise in sensors and bit quantization errors. Salt and Pepper [6] noise makes some pixles black (0) or white (1). To test the algorithms we added ρ percent of salt and pepper noise in

Fig. 9. One realization of salt and pepper noise denoising for $\rho = 0.2$. Columns from left to right: Noisy image, median filter, RB-RVM and SD-RVM. Peak Signal to Noise Ratio (PSNR) has been rounded to two decimals.

7 different images (Boat, Baboon, Barbara, Elaine, House, Lena and Peppers) and denoised them using the median filter, the RVM, the RB-RVM and the SD-RVM. The pixels were set to either black or white with equal probability.

The median filter estimates the value of each pixel by the median in a 3×3 square patch. For the RVM, RB-RVM and SD-RVM, the value of a pixel was estimated by forming a 5×5 square patch around the pixel. In the patch, the pixels were modeled as [29]

$$y_i = \beta_0 + \boldsymbol{\beta}_1^\top (\mathbf{x} - \mathbf{x}_i) + \boldsymbol{\beta}_2^\top \mathrm{vech}((\mathbf{x} - \mathbf{x}_i)(\mathbf{x} - \mathbf{x}_i)^\top) + n_i,$$

where n_i is noise, $\mathbf{x}$ is the position of the central pixel, $\mathbf{x}_i$ is the position of pixel i, $i = 1, 2, \ldots, 25$ and vech is the *half-vectorization operator* [29], i.e.

$$\mathrm{vech}\left(\begin{pmatrix} a & b \\ b & c \end{pmatrix}\right) = \begin{pmatrix} a \\ b \\ c \end{pmatrix}.$$

Given the regression parameters, the value of the central pixel is estimated as $\hat{y} = \hat{\beta}_0$. Since pixels close to the central pixel are more important, the errors are weighted by a kernel, $K(\mathbf{x}, \mathbf{x}_i)$. The estimation problem thus becomes

$$\min_{\beta_0, \boldsymbol{\beta}_1, \boldsymbol{\beta}_2} \sum_{i=1}^{P} \Big| y_i - \beta_0 - \boldsymbol{\beta}_1^\top (\mathbf{x} - \mathbf{x}_i) - \boldsymbol{\beta}_2^\top \mathrm{vech}((\mathbf{x} - \mathbf{x}_i)(\mathbf{x} - \mathbf{x}_i)^\top) \Big|^2 K(\mathbf{x}, \mathbf{x}_i),$$

where we used the kernel

$$K(\mathbf{x}, \mathbf{x}_i) = \exp\left(-||\mathbf{x} - \mathbf{x}_i||_2^2 / r^2\right) \left(1 + \mathbf{x}^\top \mathbf{x}_i\right)^p,$$

Table 2. Mean cputime (in seconds) for denoising images corrupted by salt and pepper noise averaged over 7 images.

Algorithm	Mean cputime
Median filter	5
RVM	925
RB-RVM	1154
SD-RVM	788

the kernel is a composition of a Gaussian and polynomial kernel [6, 29]. In the simulation we used $r = 2.1$ and $p = 1$ as in [6]. To avoid overfitting, it is beneficial to promote sparsity in $[\beta_0, \boldsymbol{\beta}_1^\top, \boldsymbol{\beta}_2^\top]^\top$ [30, 31, 32].

We compared the algorithms by computing the Peak Signal to Noise Ratio (PSNR)

$$\text{PSNR} = -10 \cdot \log_{10}\left(\frac{E[||\mathbf{X} - \hat{\mathbf{X}}||_F^2]}{E[\max_{i,j}|X_{ij}|^2](pq)}\right),$$

where the size of the image is $p \times q$ and the expectation is taken over the different images and realizations of the noise. All images in the simulation were of size $p = q = 256$ with $\max_{i,j}|X_{ij}|^2 = 1$. Figure 9 shows one realization of the problem, where SD-RVM gives lower PSNR than the median filter and RB-RVM. In the simulations we varied ρ and used 10 noise realizations for each image. The result is shown in figure 8.

We found that the median filter performed best for $\rho \leq 0.1$ while SD-RVM outperformed RB-RVM for all values of ρ and also the median filter for $\rho \geq 0.2$. The gain in using SD-RVM over the median filter was significant for the images Boat, Elaine, Lena, House and Peppers (for $\rho \geq 0.2$). The mean cputime of SD-RVM was 68% of the mean cputime of RB-RVM (see Table 4.4) while the median filter was by far the fastest method.

5. CONCLUSION

In this paper we introduced the combined Sparse and Dense noise Revelance Vector Machine (SD-RVM) which is robust to sparse and dense additive noise. SD-RVM was shown to be equivalent to the minimization of a non-symmetric sparsity promoting cost function. Through simulations, SD-RVM was shown to empirically perform better than the standard RVM and the robust RB-RVM.

6. APPENDIX: DERIVATION OF UPDATE EQUATIONS

Here we derive the update equations for $\hat{\mathbf{x}}$, γ_i and β_j for the SD-RVM in Sections 2.1, 3.1 and 3.2.

For fixed precisions $\boldsymbol{\gamma}$ and $\boldsymbol{\beta}$, the Maximum A Posteriori (MAP) estimate of $\mathbf{x}$ becomes

$$\begin{aligned}\hat{\mathbf{x}} &= \arg\max_{\mathbf{x}} \ \log p(\mathbf{y}, \mathbf{x}|\boldsymbol{\gamma}, \boldsymbol{\beta}) \\ &= \arg\min_{\mathbf{x}} \ (\mathbf{y} - \mathbf{A}\mathbf{x})^\top \mathbf{B}(\mathbf{y} - \mathbf{A}\mathbf{x}) + \mathbf{x}^\top \boldsymbol{\Gamma}\mathbf{x} \\ &= \boldsymbol{\Sigma}\mathbf{A}^\top \mathbf{B}\mathbf{y},\end{aligned}$$

where $\boldsymbol{\Sigma} = (\boldsymbol{\Gamma} + \mathbf{A}^\top \mathbf{B}\mathbf{A})^{-1}$. The form of the MAP estimate is the same for all models considered in this paper.

6.1. Derivation of (11) and (15)

Proof of (11). Since

$$\mathbf{B}^{-1} + \mathbf{A}\boldsymbol{\Gamma}^{-\mathbf{1}}\mathbf{A}^\top = \mathbf{B}^{-1} + \sum_{i=1}^{n} \gamma_i^{-1}\mathbf{a}_i\mathbf{a}_i^\top,$$

where $\mathbf{a}_i$ is the i'th column vector of $\mathbf{A}$ we find that

$$\begin{aligned}&\frac{\partial}{\partial \gamma_i}\left(\mathbf{y}^\top(\mathbf{B}^{-1} + \mathbf{A}\boldsymbol{\Gamma}^{-1}\mathbf{A}^\top)^{-1}\mathbf{y}\right) \\ &= \gamma_i^{-2}\left(\mathbf{a}_i^\top(\mathbf{B}^{-1} + \mathbf{A}\boldsymbol{\Gamma}^{-1}\mathbf{A}^\top)^{-1}\mathbf{y}\right)^2.\end{aligned}$$

Using that

$$\begin{aligned}&\boldsymbol{\Gamma}^{-1}\mathbf{A}^\top(\mathbf{B}^{-1} + \mathbf{A}\boldsymbol{\Gamma}^{-1}\mathbf{A}^\top)^{-1}\mathbf{y} \\ &= \boldsymbol{\Gamma}^{-1}\mathbf{A}^\top\left(\mathbf{B} - \mathbf{B}\mathbf{A}(\boldsymbol{\Gamma} + \mathbf{A}^\top\mathbf{B}\mathbf{A})^{-1}\mathbf{A}^\top\mathbf{B}\right)\mathbf{y} \\ &= \boldsymbol{\Gamma}^{-1}\mathbf{A}^\top\mathbf{B}\mathbf{y} - \boldsymbol{\Gamma}^{-1}\underbrace{\mathbf{A}^\top\mathbf{B}\mathbf{A}}_{=\boldsymbol{\Sigma}^{-1}-\boldsymbol{\Gamma}}\underbrace{(\boldsymbol{\Gamma} + \mathbf{A}^\top\mathbf{B}\mathbf{A})^{-1}}_{=\boldsymbol{\Sigma}}\mathbf{A}^\top\mathbf{B}\mathbf{y} \\ &= \boldsymbol{\Sigma}\mathbf{A}^\top\mathbf{B}\mathbf{y} = \hat{\mathbf{x}},\end{aligned} \tag{27}$$

we get that

$$\frac{\partial}{\partial \gamma_i}\left(\mathbf{y}^\top(\mathbf{B}^{-1} + \mathbf{A}\boldsymbol{\Gamma}^{-1}\mathbf{A}^\top)^{-1}\mathbf{y}\right) = \gamma_i^{-2}(\gamma_i\hat{x}_i)^2 = \hat{x}_i^2.$$

□

Proof of (15). Since

$$\begin{aligned}&\mathbf{y}^\top(\mathbf{B}^{-1} + \mathbf{A}\boldsymbol{\Gamma}^{-\mathbf{1}}\mathbf{A}^\top)^{-1}\mathbf{y} \\ &= \mathbf{y}^\top\mathbf{B}\mathbf{y} - \mathbf{y}\mathbf{B}\mathbf{A}\underbrace{(\boldsymbol{\Gamma} + \mathbf{A}^\top\mathbf{B}\mathbf{A})^{-1}}_{=\boldsymbol{\Sigma}}\mathbf{A}^\top\mathbf{B}\mathbf{y},\end{aligned}$$

we get that

$$\begin{aligned}&\frac{\partial}{\partial \beta_j}\left(\mathbf{y}^\top(\mathbf{B}^{-1} + \mathbf{A}\boldsymbol{\Gamma}^{-1}\mathbf{A}^\top)^{-1}\mathbf{y}\right) \\ &= y_j^2 - 2y_j\mathbf{A}_{j,:}\boldsymbol{\Sigma}\mathbf{A}^\top\mathbf{B}\mathbf{y} + \mathbf{y}^\top\mathbf{B}\mathbf{A}\boldsymbol{\Sigma}\mathbf{A}_{j,:}^\top\mathbf{A}_{j,:}\boldsymbol{\Sigma}\mathbf{A}^\top\mathbf{B}\mathbf{y} \\ &= y_j^2 - 2y_j\mathbf{A}_{j,:}\hat{\mathbf{x}} + (\mathbf{A}_{j,:}\mathbf{x})^2 = [\mathbf{y} - \mathbf{A}\hat{\mathbf{x}}]_j^2.\end{aligned}$$

□

6.2. Known block structure

Let $\mathbf{\Gamma}$ and $\mathbf{B}$ be diagonal matrices with

$$[\mathbf{\Gamma}]_{kk} = \gamma_i, \text{ if } k \in I_i, \qquad [\mathbf{B}]_{ll} = \beta_j, \text{ if } l \in J_j,$$

and zero otherwise.

To update the precisions we maximize the marginal distribution

$$p(\mathbf{y}, \boldsymbol{\gamma}, \mathbf{B}) = p(\mathbf{y}|\boldsymbol{\gamma}, \boldsymbol{\beta})p(\boldsymbol{\gamma})p(\boldsymbol{\beta}),$$

with respect to $\boldsymbol{\gamma}$ and $\boldsymbol{\beta}$, where $p(\boldsymbol{\gamma})$ and $p(\boldsymbol{\beta})$ is as in (3) and (19). The log-likelihood of the parameters is

$$\begin{aligned}\mathcal{L} &= \text{const.} - \frac{1}{2}\log\det(\mathbf{B}^{-1} + \mathbf{A}\mathbf{\Gamma}^{-1}\mathbf{A}^\top) \\ &- \frac{1}{2}\mathbf{y}^\top(\mathbf{B}^{-1} + \mathbf{A}\mathbf{\Gamma}^{-1}\mathbf{A}^\top)^{-1}\mathbf{y} \\ &+ \sum_{i=1}^{p}(a\log\gamma_i - b\gamma_i) + \sum_{j=1}^{q}(c\log\beta_j - d\beta_j).\end{aligned}$$

Using (12) we get that $\mathcal{L}$ is maximized when

$$\begin{aligned}\frac{\partial\mathcal{L}}{\partial\gamma_i} &= -\frac{1}{2}\text{tr}(\mathbf{\Sigma}_{I_i}) + \frac{n_i}{2\gamma_i} + \frac{a}{\gamma_i} - b \\ &- \frac{1}{2\gamma_i^2}||\mathbf{A}_{I_i}^\top(\mathbf{B}^{-1} + \mathbf{A}\mathbf{\Gamma}^{-1}\mathbf{A}^\top)^{-1}\mathbf{y}||_2^2 = 0, \qquad (28)\end{aligned}$$

where $\mathbf{\Sigma}_{I_i} \in \mathbb{R}^{n_i \times n_i}$ is the submatrix of $\mathbf{\Sigma}$ consisting of the columns and rows in I_i. Further, using (27) we get that

$$\mathbf{A}_{I_i}^\top(\mathbf{B}^{-1} + \mathbf{A}\mathbf{\Gamma}^{-1}\mathbf{A}^\top)^{-1}\mathbf{y} = \gamma_i\hat{\mathbf{x}}_{I_i}. \qquad (29)$$

Thus, (28) is fulfilled when

$$-\frac{1}{2}\text{tr}(\mathbf{\Sigma}_{I_i}) + \frac{n_i}{2\gamma_i} + \frac{a}{\gamma_i} - b - \frac{1}{2}||\hat{\mathbf{x}}_{I_i}||_2^2 = 0.$$

As before, instead of solving for γ_i we rewrite the equation as

$$n_i - \gamma_i\text{tr}(\mathbf{\Sigma}_{I_i}) + 2a - (||\hat{\mathbf{x}}_{I_i}||_2^2 + 2b)\gamma_i^{new} = 0. \qquad (30)$$

Solving (30) for γ_i^{new} gives us the update equation (20).

To find the update equation for β_j we use that

$$\begin{aligned}&\frac{\partial}{\partial\beta_j}\left[\mathbf{y}^\top(\mathbf{B}^{-1} + \mathbf{A}\mathbf{\Gamma}^{-1}\mathbf{A}^\top)^{-1}\mathbf{y}\right] = ||\mathbf{y}_{J_j}||_2^2 \\ &- 2\mathbf{y}_{J_j}^\top\mathbf{A}_{J_j,:}\mathbf{\Sigma}\mathbf{A}^\top\mathbf{B}\mathbf{y} + \mathbf{y}^\top\mathbf{B}\mathbf{A}\mathbf{\Sigma}\mathbf{A}_{J_j,:}^\top\mathbf{A}_{J_j,:}\mathbf{\Sigma}\mathbf{A}^\top\mathbf{B}\mathbf{y} \\ &= ||(\mathbf{y} - \mathbf{A}\hat{\mathbf{x}})_{J_j}||_2^2,\end{aligned}$$

where $\mathbf{A}_{J_j,:}$ consists of the row vectors of $\mathbf{A}$ which row number belongs to J_j. We get that

$$\begin{aligned}\frac{\partial\mathcal{L}}{\partial\beta_j} &= -\frac{1}{2}\text{tr}(\mathbf{\Sigma}\mathbf{A}_{J_j,:}^\top\mathbf{A}_{J_j,:}) + \frac{m_j}{2\beta_j} \\ &- \frac{1}{2}||(\mathbf{y} - \mathbf{A}\hat{\mathbf{x}})_{J_j}||_2^2 + \frac{c}{\beta_j} - d = 0.\end{aligned}$$

Rewriting the equation as

$$\begin{aligned}&1 - \beta_j\text{tr}(\mathbf{A}_{J_j,:}\mathbf{\Sigma}\mathbf{A}_{J_j,:}^\top) + 2c \\ &- (||(\mathbf{y} - \mathbf{A}\hat{\mathbf{x}})_{J_j}||_2^2 + 2d)\beta_j^{new} = 0,\end{aligned}$$

and using that $\text{tr}(\mathbf{A}_{J_j,:}\mathbf{\Sigma}\mathbf{A}_{J_j,:}^\top) = \text{tr}([\mathbf{A}\mathbf{\Sigma}\mathbf{A}^\top]_{J_j})$ gives us the update equation (21).

6.3. Unknown block structure

When the block structure is unknown, we use the over-parametrized model in section 3.2. The log-likelihood of the parameters is

$$\begin{aligned}\mathcal{L} &= \log p(\mathbf{y}|\boldsymbol{\gamma}, \boldsymbol{\beta})p(\boldsymbol{\gamma})p(\boldsymbol{\beta}) \\ &= \text{const.} - \frac{1}{2}\log\det(\mathbf{B}^{-1} + \mathbf{A}\mathbf{\Gamma}^{-1}\mathbf{A}^\top) \qquad (31) \\ &- \frac{1}{2}\mathbf{y}^\top(\mathbf{B}^{-1} + \mathbf{A}\mathbf{\Gamma}^{-1}\mathbf{A}^\top)^{-1}\mathbf{y} \\ &+ \sum_{i=1}^{p}(a\log\tilde{\gamma}_i - b\tilde{\gamma}_i) + \sum_{j=1}^{q}(c\log\tilde{\beta}_j - d\tilde{\beta}_j).\end{aligned}$$

We search to maximize (31) with respect to the underlying variables $\tilde{\gamma}_k$ and $\tilde{\beta}_l$. Using that $\frac{\partial\gamma_i^{-1}}{\partial\tilde{\gamma}_k} = -\tilde{\gamma}_k^{-2}$, $\frac{\partial\gamma_i}{\partial\tilde{\gamma}_k} = \gamma_i^2\tilde{\gamma}_k^{-2}$, when $i \in I_k$ and zero otherwise, (12) and (27) we find that $\mathcal{L}$ is maximized when

$$\begin{aligned}\frac{\partial\mathcal{L}}{\partial\tilde{\gamma}_k} &= -\frac{1}{2\tilde{\gamma}_k^2}\text{tr}(\mathbf{\Sigma}\mathbf{\Gamma}_k^2) + \frac{1}{2\tilde{\gamma}_k^2}\text{tr}(\mathbf{\Gamma}_k) \\ &- \frac{1}{2\tilde{\gamma}_k^2}\hat{\mathbf{x}}^\top\mathbf{\Gamma}_k^2\hat{\mathbf{x}} + \frac{a}{\tilde{\gamma}_k} - b = 0, \qquad (32)\end{aligned}$$

By rewriting (32) as

$$\begin{aligned}&\frac{1}{\tilde{\gamma}_k}\text{tr}(\mathbf{\Gamma}_k) - \frac{1}{\tilde{\gamma}_k}\text{tr}(\mathbf{\Gamma}_k\mathbf{\Sigma}\mathbf{\Gamma}_k) + 2a \\ &- \left(\frac{1}{\tilde{\gamma}_k^2}||\mathbf{\Gamma}_k\hat{\mathbf{x}}||_2^2 + 2b\right)\tilde{\gamma}_k^{new} = 0. \qquad (33)\end{aligned}$$

Solving (33) for $\tilde{\gamma}_k^{new}$ gives us the update equation (23).

For the noise precisions, we similarly find that

$$\begin{aligned}\frac{\partial\mathcal{L}}{\partial\tilde{\beta}_l} &= -\frac{1}{2\tilde{\beta}_l^2}\text{tr}(\mathbf{B}_l\mathbf{A}^\top\mathbf{\Sigma}\mathbf{A}\mathbf{B}_l) + \frac{1}{2\tilde{\beta}_l^2}\text{tr}(\mathbf{B}_l) \\ &- \frac{1}{2\tilde{\beta}_l^2}||\mathbf{B}_l(\mathbf{y} - \mathbf{A}\hat{\mathbf{x}})||_2^2 + \frac{c}{\tilde{\beta}_l} - d = 0.\end{aligned}$$

By rewriting the expression as

$$\begin{aligned}&\frac{1}{\tilde{\beta}_l}\text{tr}(\mathbf{B}_l) - \frac{1}{\tilde{\beta}_l}\text{tr}(\mathbf{B}_l\mathbf{A}^\top\mathbf{\Sigma}\mathbf{A}\mathbf{B}_l) + 2c \\ &- \left(\frac{1}{\tilde{\beta}_l^2}||\mathbf{B}_l(\mathbf{y} - \mathbf{A}\hat{\mathbf{x}})||_2^2 + 2d\right)\tilde{\beta}_l^{new} = 0,\end{aligned}$$

we find the update equation (24).

We see that the form of update equations depends on how the equations are rewritten. The form used here has the advantage of reducing to (21) when the underlying blocks are disjoint.

REFERENCES

[1] M. Tipping, *The relevance vector machine*, NIPS, 1999, pp. 652-658.

[2] C. Bishop, *Pattern Recognition and Machine Learning*, Springer-Verlag New York, Inc. Secaucus, NJ, USA, 2006.

[3] D.P. Wipf and B.D. Rao, *Sparse Bayesian learning for basis selection*, IEEE Transactions on Signal Processing, vol.52, no.8, pp.2153 - 2164, Aug. 2004.

[4] D. Wipf, J. Palmer and B.D. Rao, *Perspectives on sparse Bayesian learning*, Advances in neural information processing systems, vol. 16, pp. 249 - 256, 2004.

[5] S. Ji, Y. Xue and L. Carin, *Bayesian compressive sensing*, IEEE Transactions on Signal Processing, vol. 56, no. 6, pp. 2346-2356, 2008.

[6] K. Mitra, A. Veeraraghavan and R. Chellappa, *Robust RVM regression using sparse outlier model*, 2012 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2012, pp. 1887-1894.

[7] J. Laska, M. Davenport and R. Baraniuk, *Exact signal recovery from sparsely corrupted measurements through the pursuit of justice*, Proceedings of the 43rd Asimolar conference on Signals, systems and computers, Piscataway, NJ, USA, 2009, pp. 1556-1560, IEEE Press.

[8] Y. Jin and B.D. Rao, *Algorithms for robust linear regression by exploiting the connection to sparse signal recovery*, IEEE International Conference on Acoustics Speech and Signal Processing (ICASSP), 2010, pp.3830 - 3833, 14-19 March 2010.

[9] M. Vehkapera, Y. Kabashima and S. Chatterjee, *Statistical mechanics approach to sparse noise denoising*, Proceedings of the 21st European Signal Processing Conference (EUSIPCO), 2013, pp. 1-5, 9-13 September 2013.

[10] A. Cherian, S. Sra and N. Papanikolopoulos, *Denoising sparse noise via online dictionary learning*, IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 2011, pp. 2060 - 2063, 22-27 May 2011.

[11] R. Giri and B.D. Rao, *Block sparse excitation based all-pole modeling of speech*, IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 2014, pp. 3754 - 3758, 4-9 May 2014.

[12] D. Giacobello, M.G. Christensen, M.N. Murthi, S.H. Jensen and M. Moonen, *Sparse Linear Prediction and Its Applications to Speech Processing*, IEEE Transactions on Audio, Speech and Language Processing, vol.20, no.5, pp.1644 - 1657, July 2012.

[13] V. Kekatos and G.B. Giannakis, *From Sparse Signals to Sparse Residuals for Robust Sensing*, IEEE Transactions on Signal Processing, vol.59, no.7, pp. 3355 - 3368, July 2011.

[14] R.E. Carrillo, K.E. Barner and T.C. Aysal, *Robust Sampling and Reconstruction Methods for Sparse Signals in the Presence of Impulsive Noise*, IEEE Journal of Selected Topics in Signal Processing, vol.4, no.2, pp. 392 - 408, April 2010.

[15] J. Wright and Y. Ma, *Robust face recognition via sparse representation*, IEEE Transactions on Information Theory, vol. 56, no. 7, pp. 3540-3560, July 2010.

[16] Z. Zhang and B.D. Rao, *Extension of sbl algorithms for the recovery of block sparse signals with intra-block correlation*, IEEE Transactions on Signal Processing, vol. 61, no. 8, pp. 2009-2015, April 2013.

[17] Z. Zhang and B.D. Rao, *Sparse Signal Recovery With Temporally Correlated Source Vectors Using Sparse Bayesian Learning*, IEEE Journal of Selected Topics in Signal Processing, vol.5, no.5, pp.912,926, Sept. 2011.

[18] S. Chen, D. Donoho and M. Saunders, *Atomic decomposition by basis pursuit*, SIAM Rev., vol. 43, no. 1, pp. 129-159, January 2001.

[19] S. Chatterjee, D. Sundman, M. Vehkapera, M. Skoglund, *Projection-Based and Look-Ahead Strategies for Atom Selection*, IEEE Transactions on Signal Processing, vol.60, no.2, pp.634 - 647, Feb. 2012.

[20] D. Zachariah, S. Chatterjee, M. Jansson, *Dynamic Iterative Pursuit*, IEEE Transactions on Signal Processing, vol.60, no.9, pp.4967 - 4972, Sept. 2012.

[21] D. Harville, *Matrix algebra from a statistician's perspective*, Springer, 2008.

[22] D. MacKay, *Bayesian interpolation*, Neural Computation, vol. 4, pp. 415-447, 1991.

[23] C. Rojas, D. Katselis and H. Hjalmarsson, *A note on the spice method*, IEEE transactions on signal processing, vol 61. no. 18, pp. 4545-4551, 2013.

[24] L. Trefethen and D. Bau, *Numerical linear algebra*, Society for Industrial and Applied Mathematics, 1997.

[25] K. Bache and M. Lichman, *UCI machine learning repository*, 2013.

[26] M. Grant and S. Boyd, *CVX: Matlab Software for Disciplined Convex Programming, version 2.1*, http://cvxr.com/cvx, March 2014.

[27] E. Candes, J. Romberg and T. Tao, *Stable signal recovery from incomplete and inaccurate measurements*,

Communications on Pure and Applied Mathematics, vol. 59, no. 8, pp. 1207-1223, 2006.

[28] R.G. Baraniuk, V. Cevher, M.F. Duarte, C. Hegde, *Model-Based Compressive Sensing*, IEEE Transactions on Information Theory, vol.56, no.4, pp. 1982 - 2001, April 2010.

[29] H. Takeda, S. Farsiu and P. Milanfar, *Robust kernel regression for restoration and reconstruction of images from sparse noisy data*, IEEE International Conference on Image Processing, 2006.

[30] M. Elad and M. Aharon, *Image Denoising Via Sparse and Redundant Representations Over Learned Dictionaries*, IEEE Transactions on Image Processing, vol.15, no.12, pp. 3736 - 3745, Dec. 2006.

[31] J. Mairal, M. Elad and G. Sapiro, *Sparse Representation for Color Image Restoration*, IEEE Transactions on Image Processing, vol.17, no.1, pp. 53 - 69, Jan. 2008.

[32] A. Bruckstein, D. Donoho and M. Elad, *From Sparse Solutions of Systems of Equations to Sparse Modeling of Signals and Images*, SIAM Review, vol.51, no.1, pp. 34 - 81, 2009.